

ABSTRACT

Methods and systems in accordance with the present invention provide a peer-to-peer replicated hierarchical data store that allows the synchronization of the contents of multiple data stores on a computer network without the use of a master data store. The synchronization of a replicated data store stored on multiple locations is provided even when there is constantly evolving set of communications partitions in the network. Each computer in the network may have its own representation of the replicated data store and may make changes to the data store independently without consulting a master authoritative data store or requiring a consensus among other computers with representations of the data store. Changes to the data store may be communicated to the other computers by broadcasting messages in a specified protocol to the computers having a representation of the replicated data store. The computers receive the messages and process their local representation of the data store according to a protocol described below. As such, each computer has a representation of the replicated database that is consistent with the representations of the data store on the other computers. This allows computers to make changes to the data store even when disconnected via a network partition.